

METHOD AND APPARATUS FOR STORING CONTENT WITHIN A VIDEO ON DEMAND ENVIRONMENT

CROSS REFERENCE

- 5 This application claims benefit of United States Provisional Application No. 60/170,138, filed December 10, 1999, which is hereby incorporated by reference in its entirety.

BACKGROUND OF THE INVENTION

10 1. Field of the Invention

The present invention relates to an information distribution system such as a video-on-demand (VOD) system. More particularly, the present invention relates to a method and apparatus for efficiently utilizing "on line" and "near line" storage media to provide increased availability of content streams to users.

2. Description of the Background Art

20 VOD systems providing content encoded according to the various Moving Pictures Experts Group (MPEG) standards are known. For example, a first standard known as MPEG-1 refers to ISO/IEC standards 11172, which is incorporated herein by reference in its entirety. A second standard known as MPEG-2 refers to ISO/IEC standards 13818, which is incorporated herein by reference in its entirety. Additionally, a compressed digital video system is described in the Advanced Television Systems Committee (ATSC) digital television standard document A/53, incorporated herein by reference.

30 The most important characteristic of a video on demand system is the experience of a user receiving content via that system. That is, a user ideally selects desired content and the system provides the desired content to the user in a timely manner and at a quality level consistent with the user's expectations. Failure to meet the user expectations in terms of image

003037 404460

quality, sound quality, system latency or other factors results in a less than satisfying experience to the user.

In the case of a video on demand system utilizing a plurality of servers to provide content to users, it is seen to be desirable to sense anomalous conditions or other errors within a server presently providing content to a user and, upon determination that an error exists, migrating that user to a server not affected by the error condition. More specifically, it is seen to be desirable to provide a method and apparatus for migrating users between information server modules in a manner minimizing disruption of content streams being provided to the user.

SUMMARY OF THE INVENTION

The disadvantages heretofore associated with the prior art are overcome by the present invention of a method and apparatus for maximizing the number of content titles available within a video on demand system by storing an initial portion of each title on the primary storage device and a remaining portion of each title on a secondary storage device, wherein said secondary storage device is used to provision said primary storage device upon request of said content stream stored on said primary storage device.

BRIEF DESCRIPTION OF THE DRAWINGS

The teachings of the present invention can be readily understood by considering the following detailed description in conjunction with the accompanying drawings, in which:

FIG. 1 depicts a high level block diagram of an interactive information distribution system;

FIG. 2 depicts a block diagram of an information server suitable for use in the interactive information distribution system of FIG. 1;

FIG. 3 depicts a graphical diagram useful in understanding the present invention;

FIG. 4 depicts a graphical diagram useful in understanding the present invention; and

FIG. 5 depicts a flow diagram of a method according to the present invention and suitable for use in the interactive information distribution system of FIG. 1

To facilitate understanding, identical reference numerals have been used, where possible, to designate identical elements that are common to the figures.

DETAILED DESCRIPTION

FIG. 1 depicts a high level block diagram of an interactive information distribution system. Specifically, FIG. 1 depicts a high level block diagram of an interactive information distribution system 100 containing the present invention. The system 100 contains service provider equipment 102, a communications network 104 and subscriber equipment 106_n, where n is an integer greater than zero.

The service provider equipment 102 comprises an information server 125, a session controller 145, a transport processor 150 and, optionally, a remote storage module 128. Briefly, the session controller 145, in response to a request(s) from subscriber equipment 106, causes the requested content to be retrieved from the information server 125 and provided to the transport processor 150. The transport processor 150 combines or multiplexes the retrieved content to provide an output data stream for the requesting subscriber(s). The output data stream is conditioned for transport to the requested subscriber via a forward application transport channel (FATC) within the distribution network 104.

The information server 125 is used to store at least a portion of content such as movies, television programs and other information offerings of the interactive information distribution system 100 of FIG. 1. Additionally, the information server 125 is used to store assets such as bit map imagery, graphic overlay, control scripts and the like. The assets may comprise, for example, navigation assets that are used by a set top terminal to interactively navigate, and select for viewing, the offerings or content available from the service provider equipment 102. The information server 125, in response to a control SC produced by the session controller 145, provides content and/or asset data to the transport processor 150. In the

5

10

30

[illegible]

5
10
15
20

25

30

portion from the remote secondary storage 115R or remote storage module 128R includes sufficient bandwidth. It is also noted that the secondary storage devices 115, 115R and the remote storage module 128 may store portions or entireties of content streams.

5 In one embodiment of the invention, a primary storage device includes only an initial portion of a content stream to be provided to a user. In response to the user selecting the content stream, a secondary storage device is immediately utilized to provision the primary storage device with the remaining portion of the content stream. Thus, the primary storage
10 device is used to store the entirety of a content stream requested by a user, and such provisioning of the primary storage device is effected upon, for example, a user request for the content stream.

In alternate embodiments of the invention, the secondary storage device streams the remaining portion of content to the requesting user while, optionally, provisioning the primary storage device with the remaining portion of the content stream. In this manner, multiple users requesting the same content stream may be satisfied via a single access of the secondary storage device.

In the case where the primary storage device is provisioned by the secondary storage device, the primary storage device may be used to stream the remaining portion of content after such provisioning. In this example, the initial portion of a content stream requested by a user is provided by a primary storage device, a first portion of the remaining portion is then provided by the secondary storage device which also contemporaneously provisions the primary storage device with the remaining portion of the content stream. Upon completion of the provisioning of the primary storage device, the primary storage device is again used to provide the content stream directly to the user. Thus, transitions in servicing the user are made from the primary storage device to the secondary storage device at the end of the stored initial portion of the content, and from the secondary storage device to the primary storage device after provisioning of the primary storage device. Such transition may be effected by migrating users between primary and secondary storage devices/servers.

10

15

20

30

In one embodiment of the invention, each of the output buffers 225 supporting the server modules 220 comprise respective portions of a common memory module. That is, each of the server modules 220 stores data to a respective portion of a common memory module. In this
5 embodiment of the invention, the switch 230 utilizes a direct memory access (DMA) output link table 235 to access the appropriate portions of the common buffer memory to retrieve the contents of each of the output buffers 225.

Each server module 220 is capable of providing information to a
10 plurality of users 106. Thus, each buffer 225 associated with a server module 220 is capable of holding at least one extent of data for each of the plurality of subscribers 106 serviced by that server module 220. For example, if the first server module (220₁) is capable of serving 100 subscribers, then the buffer 225₁ associated with the first server module
15 220₁ must be capable of holding at least 200 seconds worth of information, illustratively video information and any associated audio information.

Under normal operating circumstances, each server module 220 retrieves information from its respective disk array (primary storage) for each subscriber 106 supported, and transfers the retrieved information to
20 the respective buffer (or buffer region) 225. The switch 230 accesses each buffer 225 in a round robin fashion to produce an output stream OUT comprising the multiplexed contents of each of the buffers 225₁ to 225_m. In the case of only a portion of each available information stream being stored in a respective disk array, the stored portion is transferred to the respective
25 buffer 220 while at the same time the remaining portion is retrieved from a secondary storage unit, such as local secondary storage module 115 or remote secondary storage module 128.

In the case of a server module failure, a primary or secondary storage failure such as a disk array failure, a buffer failure, or some other event
30 affecting the flow of information to subscribers via one of the server modules 220, the users on the affected server module 220 is migrated to another server module. Since the information server 125 comprises a plurality of server modules 220, the loss of one or more server modules 220 may be compensated for by migrating some or all of the users from the damaged

003340 12000

server module(s) to the remaining, functional server modules. In this manner, component failures within the information server 125 will result in a degradation of information server performance that, ideally, will not adversely impact the experience of users receiving content streams within the information distribution system 100 of FIG. 1.

Advantageously, a user may also be migrated to another server module if the other server module includes, in a respective primary storage module, the entire content stream requested by that user. Specifically, assume that a first server module is servicing a user that has requested a particular content stream. If the particular content stream is only partially stored in the primary storage associated with the first server module, then the remaining portion of the requested content must be retrieved from a secondary storage module. However, if another server module includes the entirety of the requested content within its primary storage, then the user may be migrated to the second server module. In this manner, the user content request may be satisfied without resorting to secondary storage access. Apparatus and method for effecting a user migration between server modules is described in more detail in commonly assigned U.S. Patent Application Serial No. _____, filed on _____ (Attorney Docket No. 533/241), which is incorporated herein by reference in its entirety.

To accomplish such migration of users between server modules 220, and to determine if requested content is available on other server modules 220, the server modules 220 communicate with each other via a service module bus SMBUS to determine if such a migration is possible or useful.

FIG. 3 depicts a graphical diagram useful in understanding the present invention. Specifically, FIG. 3 depicts a graphical representation of content stored by a primary storage module and a secondary storage module according to one aspect of the invention. Specifically, a primary storage module, such as the disk array 110, is used to store only a portion of each of a plurality of content streams. Each content stream comprises at least a play track and, optionally, a fast forward track and rewind track. A play track comprises a normal speed track. A fast forward track comprises a temporally decimated version of a play track which, when played by the

subscriber equipment 106, results in a fast forward or "special play" presentation of the content on the user's presentation device. A rewind track comprises a temporally decimated version of a play track which has been reversed in order (e.g., a frame-reversed FF track) such that, when
5 played by a user, results in a rewind or "reverse play" presentation of the content. It should be noted that fast forward and rewind tracks may be created to impart any desired increase in apparent play rate. However, in keeping with the traditional fast forward/rewind rates of standard analog video cassette recorders (VCRs), the exemplary embodiment uses fast
10 forward rewind tracks having a presentation rate of between seven and nine times the presentation rate of the play track.

Referring to the primary storage 110 depicted in FIG. 3, a plurality of content streams M1-MZ (310₁- 310_z) are depicted as being represented by respective initial play (P), fast forward (FF) and rewind (R) portions. Each
15 of the respective initial portions of the content streams M1-MZ (310₁ - 310_z) is sufficient to provide an initial presentation of, for example, 20 minutes to a requesting user. Additionally, a portion 320 of primary storage is reserved for the remainder of content being streamed to a requesting user. Specifically, a reserved portion 320 of primary storage is used to store the
20 remainder portion of content stream(s) requested by user(s) within the system.

Referring to the secondary storage 115/128 depicted in FIG. 3, a plurality of at least remaining portions of the content streams M1-MZ (330₁- 330_z) are depicted as being stored therein. For example, in the case of a user
25 requesting a first content stream M1 (310₁), the play track of the first content stream M1P is streamed to the requesting user. The remaining portion of the first content stream M1 is retrieved from the secondary storage module such as local secondary storage module 115 or remote secondary storage module 128 and stored in the reserved portion 320 of the
30 primary storage module 110 for subsequent streaming to the requesting user. It is important to store at least some of the remaining portion of the requested content in reserved portion 320 prior to the presentation of the terminating image frame F_{Y1}. In this manner, the stream provided to the

user may be changed from the initial content stream 310 to a remaining content stream 320 in a relatively seamless manner.

Referring to FIG. 3, a primary storage module (illustratively a disk array 110) is depicted as storing respective initial portions of each of a plurality of content streams (denoted as M1 through MZ while a secondary storage module (illustratively a magneto-optical or tape drive) is depicted as storing respective remaining portions of the plurality of content streams). While the following discussion is primarily directed to a first content stream M1, it will be understood that the discussion is equally applicable to the other content streams M2-MZ.

Specifically, a portion of a first content stream M1 is stored as a play track portion M1P, a fast forward track portion M1FF and a rewind track portion M1R. The play track portion M1P comprises all frames between a first image frame F_{01} and a terminating image frame F_{Y1} . The first image frame F_{01} comprises the first image frame of the content stream M1. The terminating image frame F_{Y1} comprises the last image frame of the initial portion of the content stream M1 stored in the primary storage device. For purposes of this discussion, it will be assumed that a 20 minute portion of each content stream M1 - MZ is stored in the primary storage device. Therefore, assuming a frame rate of 30 frames per second, the terminating image frame F_{Y1} is approximately the 36,000th frame of the first content stream M1.

The fast forward track portion M1FF of the content stream M1 comprises that portion of the fast forward track including the first image frame F_{01} and the terminating image frame F_{Y1} . Similarly, the portion of the rewind track M1R of the content stream M1 comprises that portion of the rewind track including the terminating image frame F_{Y1} and the first image frame F_{01} . In the embodiment of the invention depicted in FIG. 3, the rewind track M1R of the content stream M1 comprises, essentially, a reversed frame order version of the fast forward track M1FF of the content stream M1. However, as will be discussed in more detail below, the fast forward and rewind tracks do not need to be of the same length, nor must they be inclusive of the same content. However, to simply the discussion of

the present invention, the fast forward and rewind tracks are made symmetrical as shown in FIG. 3.

Thus, an initial (e.g., 20 minutes) portion of a first content stream M1 is represented by those portions of a corresponding play track M1P, fast forward track M1FF and rewind track M1R bounded by an initial image frame F_{01} and a terminating image frame F_{Y1} . Though not shown, it will be appreciated by those skilled in the art that at least the audio portion associated with the play track M1P is also stored in the primary storage device and presented with the play track. In addition to the audio portion, meta-data and other auxiliary data related to the content stream may also be stored on the primary storage device for utilization within the presentation of the content by a user.

In one embodiment of the invention, the fast forward track M1FF and rewind track M1R associated with the content stream do not include all of the corresponding frames associated with the entire play track M1P. Specifically, since the major use for a fast forward track, at least at the beginning of a content stream, is fast forwarding through introductory materials, trailers, credits and other content other than the feature presentation of the content stream, such non-feature presentation content is included within the fast forward track M1FF. In this manner, the storage of fast forward (and, optionally, rewind) track information associated with feature presentation is avoided.

Referring to FIG. 3, it is noted that the secondary storage module is depicted as including the entire play, fast forward, and rewind tracks for each of a plurality of content streams denoted as M1-MZ (330_1 - 330_z). Specifically, the entire play track M1P, fast forward track M1FF and rewind track M1R of the entire content stream M1 is depicted as being stored in a memory region 330_1 of the secondary storage module. Similarly, the entire play track, fast forward track and rewind track of content streams 2-Z are depicted as being stored in respective regions 330_2 - 330_z . However, in another embodiment of the invention, only the remaining portion of the content stream is stored in the secondary storage module. That is, the primary storage module 110 is used to store only the initial portion, while the secondary storage module is used to store only the remaining.

FIG. 4 depicts a graphical diagram useful in understanding the present invention. Specifically, FIG. 4 depicts a graphical representation of the play track 410, fast forward track 420 and rewind track 430 of a content stream.

5 A first or initial portion PT_p of the play track is stored on a primary storage module, while at least a remaining portion PT_s of the play track 410 is stored on a secondary storage module. A first or initial portion FFT_p of the fast forward track 420 is stored on the primary storage module, while a remaining portion FFT_s of the fast forward track is stored on the secondary storage module. A first or initial portion RT_p of the rewind track 430 is stored on a primary storage module, while a remaining portion RT_s of the rewind track 430 is stored on a secondary storage module.

10 It is assumed that the portion PT_p of the play track 410 stored on the primary storage module comprises a 20 minute play track. Referring to FIG. 4, the portion PT_p of the play track stored on the primary storage module begins at a time t_0 with an initial image frame F_0 , and terminates at a time t_y with a terminating image frame F_y . Each of the fast forward track portion FFT_p and rewind track portion RT_p begin at time t_0 and terminate at a time t_x . The fast forward track portion FFT_p begins with an image frame F_0 and terminates with an image frame F_y . By contrast, the rewind track portion RT_p begins with an image frame F_y and terminates at the image frame F_0 .

20 The time t_0 is the time at the start of each track stored in the primary storage module. The time t_y is equal to the time at which the play track portion PT_p ends. The time t_x is the time at which the fast forward track portion FFT_p and rewind track portion RT_p ends. The time t_x is related to t_y by the fast forward/rewind rate of the fast forward track 420 and rewind track 430. For example, if the FF/REW rate is equal to approximately nine times the play rate, then t_y is equal to approximately nine times t_x .
30 Therefore, if t_y is determined to be 20 minutes, then t_x is equal to approximately 2.22 minutes.

FIG. 5 depicts a flow diagram of a method according to the present invention and suitable for use in the system of FIG. 1. Specifically, FIG. 5 depicts a flow diagram of a method 500 for satisfying user requests for



5

10

20

30

At step 516 the remainder of the requested content is accessed from the secondary storage device determined in step 514. The method 500 then proceeds to step 518, where the server module 220 servicing the requesting user is provisioned with the remainder of the requested content and any auxiliary streams. That is, at step 518 the remaining portion of the requested content is stored in a primary storage device associated with the server module servicing the requesting user.

It should be noted that, as previously discussed, a secondary storage device is (typically) unable to provide real time streaming of content to a user, while a primary storage device is capable of providing real time streaming of content to a user. Thus, it is likely that the newly provisioned content on the primary storage device will begin streaming to the requesting user prior to the conclusion of the provisioning process.

Most important is that the user's experience is not significantly degraded or interrupted. Thus, as long as the secondary storage unit (115 or 128) provisions the primary storage unit (display 110) at a sufficiently high rate such that the transport processor 150 is able to continuously provide play, fast forward, rewind and other tracks to the user, then the system is working effectively. The method 500 then proceeds to step 520.

At step 520 the requesting user is transitioned to the newly provisioned remainder stream(s) as appropriate. That is, at step 520 a transition is made in the content stream supplied to the requesting user from the initial content stored on the primary storage device to the remaining content that has been provisioned onto the primary storage device. In the case of a second primary storage device including the requested content, the first primary storage device is not provisioned. Rather, the user stream is retrieved from the second primary storage device.

With the above-described method 500, the user receives a requested content stream without any visual or aural artifacts, while the provider 102 efficiently utilizes primary storage in a manner making a large number of titles available to each user. Such efficient storage utilization helps drive down system costs and improve system reliability. The method 500 then proceeds to step 522 where it is exited.

Optionally, after determining the location of the remainder of the content requested by the user at step 514, the above method 500 proceeds to step 530. At step 530 a determination is made (step 532) as to whether the remaining content is located on a primary storage module associated with another server module 220 within the information server 125. That is, at step 532 a determination is made as to whether another server module has associated with it a primary storage device including the content requested by the user. If the query at step 532 is answered negatively, then the method 500 proceeds to step 516. If the query at step 532 is answered affirmatively, then the method 500 proceeds to step 534.

At step 534 the user is migrated from the present server module to the other server module. That is, since another server module includes, in a respective primary storage means, the content requested by the user, that requesting user is migrated over to the other circuit module. Apparatus and method for effecting this migration between server modules is described in more detail in commonly assigned U.S. Patent Application Serial No. _____, filed on _____ which is incorporated herein by reference in its entirety.

After migrating the requesting user from the present server module to the server module including the requested content stored in the respective primary storage means, the method 500 proceeds to step 536 where it is exited.

In the exemplary embodiments of the invention discussed above, the illustrative portion of the play track stored within the primary storage is approximately 20 minutes. The 20 minute storage amount is determined with respect to the speed of the secondary storage devices and the possible use of a fast forward by a user to quickly traverse the entire initial content portion. For example, in the interactive information distribution system 100 of FIG. 1, a user requesting a content stream will begin to receive the content stream from the primary storage. If the user immediately depresses a "fast forward" key on an input device 138, then the set top terminal 136 associated with the user selects from the FATC, the fast forward stream corresponding to the presently received content stream. Thus, a user may rapidly traverse the entire portion of content stored in the primary storage.

0033407 2042260

In the case of a fast forward stream operating at nine times the frame rate of the play stream, a user will be able to traverse the entire 20 minutes of play track in only 2.22 minutes. Therefore, it is necessary that the remaining portion of the content be provided to the primary storage or, alternatively, the user be migrated to a server module including the requested content, prior to 2.22 minutes after beginning presentation of the requested content stream. While not the normal or likely user action, a user who has already seen most of the requested content may simply fast forward to that portion of the content which has not previously been viewed. Thus, it is important to insure that the provisioning of primary storage or migration of user be accomplished in a manner that "stays ahead" of the fast forward or play track demands put upon the system by the user. For example, in the case of a relatively slow secondary storage device, such as a sequentially accessed magnetic tape device in which content may be stored toward the middle or even end of a tape, it may be necessary to increase the size of the initial portion of the content stream stored on the primary storage device. Thus, the primary storage may be used to store anywhere from five minutes to 40 minutes of content, based upon the speed of secondary storage or other determination directed toward assuring a continued, uninterrupted stream of content to a requesting user. The amount of initial content storage is adapted to the system's storage capabilities (i.e., speed, access time, primary and secondary storage mix and the like) and to the acceptable acceptable probabilistic risk level of dropping a stream (i.e., allow only partial fast forward within the initial content and other limitations).

The portion of content rewind track stored in the primary storage must also be adapted to the ability of the system. Specifically, in the case of a user fast forwarding through the initial portion of content stored in the primary storage, that user may decide to rewind at some point. It is necessary to provide a rewind stream in which the user may move from a fast forward (or play) track to the rewind track. Thus, it is desirable to include a rewind track including those frames within the fast forward track.

It is noted that the play track, fast forward track and rewind tracks are all provided to the requesting user based upon user interaction. Each of

1. The first part of the document discusses the importance of maintaining accurate records of all transactions and the role of the accounting department in ensuring the integrity of the financial statements.

2. The second part of the document discusses the importance of maintaining accurate records of all transactions and the role of the accounting department in ensuring the integrity of the financial statements.

3. The third part of the document discusses the importance of maintaining accurate records of all transactions and the role of the accounting department in ensuring the integrity of the financial statements.

4. The fourth part of the document discusses the importance of maintaining accurate records of all transactions and the role of the accounting department in ensuring the integrity of the financial statements.

5. The fifth part of the document discusses the importance of maintaining accurate records of all transactions and the role of the accounting department in ensuring the integrity of the financial statements.

6. The sixth part of the document discusses the importance of maintaining accurate records of all transactions and the role of the accounting department in ensuring the integrity of the financial statements.

7. The seventh part of the document discusses the importance of maintaining accurate records of all transactions and the role of the accounting department in ensuring the integrity of the financial statements.

8. The eighth part of the document discusses the importance of maintaining accurate records of all transactions and the role of the accounting department in ensuring the integrity of the financial statements.

9. The ninth part of the document discusses the importance of maintaining accurate records of all transactions and the role of the accounting department in ensuring the integrity of the financial statements.

10. The tenth part of the document discusses the importance of maintaining accurate records of all transactions and the role of the accounting department in ensuring the integrity of the financial statements.

15 Although various embodiments which incorporate the teachings of the present invention have been shown and described in detail herein, those skilled in the art can readily devise many other varied embodiments that still incorporate these teachings.